ABSTRACT OF THE DISCLOSURE

Techniques are provided for (1) extending SQL to support direct invocation of frequent itemset operations, (2) improving the performance of frequent itemset operations by clustering itemset combinations to more efficiently use previously produced results, and (3) making on-the-fly selection of the occurrence counting technique to use during each phase of a multiple phase frequent itemset operation. When directly invoked in an SQL statement, a frequent itemset operation may receive input from results of operations specified in the SQL statement, and provide its results directly to other operations specified in the SQL statement. By clustering itemset combinations, resources may be used more efficiently by retaining intermediate information as long as it is useful, and then discarding it to free up volatile memory. Dynamically selecting an occurrence counting technique allows a single frequent itemset operation to change the occurrence counting technique that it is using midstream, based on cost considerations and/or environmental conditions.